

SEARCH SYSTEM USING USER BEHAVIOR DATA

FIELD OF THE INVENTION

[0001] This invention relates in general to the field of data search and retrieval. More particularly, this invention relates to the collection and use of user behavior data for data search and retrieval.

BACKGROUND OF THE INVENTION

[0002] Data on one or more computer systems may contain data useful for a user. However, the data may be too large for the user to find the data by direct examination of the data. Additionally, some parts of the data repository may contain information that is not accessible to the user. In many cases, in order to allow the user useful access to the data, a search mechanism is provided. The search mechanism allows a user to issue a search request (also termed a search query). The results are then returned for the user.

[0003] For example, a web-based search engine is a search mechanism which may be used to provide search access to information via a web-based search. The information may be a specific data repository, such as a database or other data collection. The information may also be an agglomeration of a number of different data repositories. Such a search engine may provide search access to information available from different information providers over a network, such as the Internet.

[0004] In a typical usage of a web search engine, the user enters a query, which is a set of search terms related to the type of content for which the user is searching. The query is transmitted to the search engine, which attempts to locate “hits” – i.e., content that is available on the Internet and that relates to the terms contained in the query. Generally, the search engine either has a database of web pages that are known to exist, or communicates with external “providers” who maintain such databases; the query is “scored” against items in these databases to identify the web pages that best match the query. A list of results is then generated, and these results are returned to the user’s computer for display by the user’s web browser.

[0005] Typically, the databases contain information such as: the Uniform Resource Locators (URLs) of web pages, the titles of the pages, descriptions of the pages, and possibly other textual information about the web pages. The user then reads the results and attempts to determine, based on the text contained in the results, whether the results correspond to what the user is looking for. Users may then attempt to retrieve the entire page correlating to a search result. In other contexts, search engines present results summarizing the pieces of data which may possibly be useful for a user.

[0006] The utility of the search engine is directly correlated to the quality of the results provided. In the best case, the most results presented to the user are presented in order of utility to the user on the result page.

[0007] Because the quality of the results is subjective, in order to determine what the quality of results are, the user’s satisfaction must be determined. For example, a user can be allowed to use a search engine for a period of time and then orally interviewed by an interviewer to determine the user’s satisfaction.

[0008] In the prior art, quality of individual web pages has been measured by obtaining explicit feedback from a user. At least one prior art web browser has attempted to obtain such explicit feedback from a user. This browser is described in a paper entitled “Inferring User Interest” by Mark Claypool, David Brown, Phong Le, Makoto Waseda in IEEE Internet Computing 5(6): 32-39 (2001). In this browser, different pages are displayed by the browser. Whenever the page being displayed by the browser is changed, a user evaluation of the page is requested from the user. User evaluations for a given page are collected, to determine whether users find that page valuable. In this browser, some implicit feedback is also maintained regarding each page, including data regarding the time spent on the page, mouse movements, mouse clicks, and scrolling time.

[0009] While this technique does gather user feedback, it has limited utility in situations in which users may have different needs for a page. For example, a user who is looking for information about books written by Douglas Adams may evaluate a page on his book The Hitchhiker's Guide to the Galaxy and give a high score for utility. However, another user who is looking for information on books about traveling cheaply may evaluate the same page and give it a low score. Thus the technique described will have limited utility in the wide variety of situations in which different users may have different needs, or even where a single user may have different needs for information at different times. In other words, the usefulness of this technique is limited because evaluation of each page is completely independent of the context in which the user arrived at the page.

[0010] Thus, this technique is not useful for evaluating the quality of a search engine. In general, this technique is not useful for evaluations which are context-based, but only for evaluating the quality of individual data items, independent of the context in which a user arrived at the data items.

[0011] In view of the foregoing, there is a need for a system and method that overcomes the drawbacks of the prior art.

SUMMARY OF THE INVENTION

[0012] User feedback (explicit and/or implicit) on a search performed on a search mechanism is collected along with the context for that feedback. This allows context-dependent evaluations of the quality of search mechanisms.

[0013] Context-based user feedback is gathered regarding searches performed on a search mechanism. The search mechanism is monitored for user behavior data regarding an interaction of a user with the search mechanism. The response data provided by the search mechanism is also monitored. Context data (describing the search) and user feedback data (the user's feedback on the search – either explicit or implicit) are determined. This can be used, for example, to evaluate a search mechanism or to check a relevance model.

[0014] Other features of the invention are described below.

BRIEF DESCRIPTION OF THE DRAWINGS

[0015] The foregoing summary, as well as the following detailed description of preferred embodiments, is better understood when read in conjunction with the appended drawings. For the

purpose of illustrating the invention, there is shown in the drawings example constructions of the invention; however, the invention is not limited to the specific methods and instrumentalities disclosed. In the drawings:

[0016] FIG. 1 is a block diagram of an example computing environment in which aspects of the invention may be implemented;

[0017] FIG. 2 is a block diagram of an example architecture according to one embodiment of the invention;

[0018] FIG. 3 is a diagram of a state machine according to one embodiment of the invention; and

[0019] FIG. 4 a flow diagram showing the process of gathering context-based feedback according to one embodiment of the invention.

DETAILED DESCRIPTION OF ILLUSTRATIVE EMBODIMENTS

Overview

[0020] User feedback (explicit and/or implicit) is collected along with the context for that feedback. This allows context-dependent evaluations of the quality of search mechanisms. If one search mechanism produces results which garner better feedback than a second search mechanism, it can be determined that the first search mechanism is more successful at finding results for a user.

[0021] Additionally, it can be used for other purposes. For example, a relevance model can be used to measure the effectiveness of a search system. A relevance model is a probabilistic model which attempts to model the usefulness of results obtained from a search. However, for such a model to itself be useful, it must be accurate. Feedback from actual users can be used to validate a relevance model by showing whether there is a correlation between the predicted utility of results (from the relevance model) and the actual utility of results (from user feedback).

[0022] Other uses may also be found for context-based user feedback monitoring. For example, the success of a sponsored link may be evaluated. If, for example, a sponsored link is returned as a search result, the determination of whether it is considered a high-quality search result may be useful for the sponsor of the link or for the search mechanism provider. Additional uses are also contemplated.

[0023] In one embodiment, context-based search mechanism user feedback is captured by tracking an event stream and following a state machine in order to determine the state of the user's

interaction with the browser. This allows queries to be presented to the user in order to (1) track what state the interaction is in and (2) request explicit feedback from the user. Implicit feedback, such as the amount of time spent on a page, among other implicit feedback items, is also tracked. This feedback information allows a determination to be made of whether a search result was useful in the context of the original search. Once this data has been collected, it can be analyzed to determine the quality of the search mechanism used. It can also be used to test a relevance model and see if the predictions generated by the relevance model on the quality of the results generated match the actual quality of the results as shown in implicit and/or explicit user feedback.

[0024] Thus, according to one embodiment of the invention, a means is provided for collecting implicit end-user behavior within search systems and correlating that behavior to user reported satisfaction with a particular search result and with the overall result set.

Exemplary Computing Environment

[0025] FIG. 1 shows an exemplary computing environment in which aspects of the invention may be implemented. The computing system environment 100 is only one example of a suitable computing environment and is not intended to suggest any limitation as to the scope of use or functionality of the invention. Neither should the computing environment 100 be interpreted as having any dependency or requirement relating to any one or combination of components illustrated in the exemplary operating environment 100.

[0026] The invention is operational with numerous other general purpose or special purpose computing system environments or configurations. Examples of well known computing systems, environments, and/or configurations that may be suitable for use with the invention include, but are not limited to, personal computers, server computers, hand-held or laptop devices, multiprocessor systems, microprocessor-based systems, set top boxes, programmable consumer electronics, network PCs, minicomputers, mainframe computers, embedded systems, distributed computing environments that include any of the above systems or devices, and the like.

[0027] The invention may be described in the general context of computer-executable instructions, such as program modules, being executed by a computer. Generally, program modules include routines, programs, objects, components, data structures, etc. that perform particular tasks or implement particular abstract data types. The invention may also be practiced in distributed computing environments where tasks are performed by remote processing devices that are linked through a communications network or other data transmission medium. In a distributed computing

environment, program modules and other data may be located in both local and remote computer storage media including memory storage devices.

[0028] With reference to FIG. 1, an exemplary system for implementing the invention includes a general purpose computing device in the form of a computer 110. Components of computer 110 may include, but are not limited to, a processing unit 120, a system memory 130, and a system bus 121 that couples various system components including the system memory to the processing unit 120. The processing unit 120 may represent multiple logical processing units such as those supported on a multi-threaded processor. The system bus 121 may be any of several types of bus structures including a memory bus or memory controller, a peripheral bus, and a local bus using any of a variety of bus architectures. By way of example, and not limitation, such architectures include Industry Standard Architecture (ISA) bus, Micro Channel Architecture (MCA) bus, Enhanced ISA (EISA) bus, Video Electronics Standards Association (VESA) local bus, and Peripheral Component Interconnect (PCI) bus (also known as Mezzanine bus). The system bus 121 may also be implemented as a point-to-point connection, switching fabric, or the like, among the communicating devices.

[0029] Computer 110 typically includes a variety of computer readable media. Computer readable media can be any available media that can be accessed by computer 110 and includes both volatile and nonvolatile media, removable and non-removable media. By way of example, and not limitation, computer readable media may comprise computer storage media and communication media. Computer storage media includes both volatile and nonvolatile, removable and non-removable media implemented in any method or technology for storage of information such as computer readable instructions, data structures, program modules or other data. Computer storage media includes, but is not limited to, RAM, ROM, EEPROM, flash memory or other memory technology, CDROM, digital versatile disks (DVD) or other optical disk storage, magnetic cassettes, magnetic tape, magnetic disk storage or other magnetic storage devices, or any other medium which can be used to store the desired information and which can be accessed by computer 110. Communication media typically embodies computer readable instructions, data structures, program modules or other data in a modulated data signal such as a carrier wave or other transport mechanism and includes any information delivery media. The term “modulated data signal” means a signal that has one or more of its characteristics set or changed in such a manner as to encode information in the signal. By way of example, and not limitation, communication media includes wired media such as a wired network or direct-wired connection, and wireless media such as

acoustic, RF, infrared and other wireless media. Combinations of any of the above should also be included within the scope of computer readable media.

[0030] The system memory 130 includes computer storage media in the form of volatile and/or nonvolatile memory such as read only memory (ROM) 131 and random access memory (RAM) 132. A basic input/output system 133 (BIOS), containing the basic routines that help to transfer information between elements within computer 110, such as during start-up, is typically stored in ROM 131. RAM 132 typically contains data and/or program modules that are immediately accessible to and/or presently being operated on by processing unit 120. By way of example, and not limitation, FIG. 1 illustrates operating system 134, application programs 135, other program modules 136, and program data 137.

[0031] The computer 110 may also include other removable/non-removable, volatile/nonvolatile computer storage media. By way of example only, FIG. 1 illustrates a hard disk drive 140 that reads from or writes to non-removable, nonvolatile magnetic media, a magnetic disk drive 151 that reads from or writes to a removable, nonvolatile magnetic disk 152, and an optical disk drive 155 that reads from or writes to a removable, nonvolatile optical disk 156, such as a CD ROM or other optical media. Other removable/non-removable, volatile/nonvolatile computer storage media that can be used in the exemplary operating environment include, but are not limited to, magnetic tape cassettes, flash memory cards, digital versatile disks, digital video tape, solid state RAM, solid state ROM, and the like. The hard disk drive 141 is typically connected to the system bus 121 through a non-removable memory interface such as interface 140, and magnetic disk drive 151 and optical disk drive 155 are typically connected to the system bus 121 by a removable memory interface, such as interface 150.

[0032] The drives and their associated computer storage media discussed above and illustrated in FIG. 1, provide storage of computer readable instructions, data structures, program modules and other data for the computer 110. In FIG. 1, for example, hard disk drive 141 is illustrated as storing operating system 144, application programs 145, other program modules 146, and program data 147. Note that these components can either be the same as or different from operating system 134, application programs 135, other program modules 136, and program data 137. Operating system 144, application programs 145, other program modules 146, and program data 147 are given different numbers here to illustrate that, at a minimum, they are different copies. A user may enter commands and information into the computer 20 through input devices such as a keyboard 162 and pointing device 161, commonly referred to as a mouse, trackball or touch pad.

Other input devices (not shown) may include a microphone, joystick, game pad, satellite dish, scanner, or the like. These and other input devices are often connected to the processing unit 120 through a user input interface 160 that is coupled to the system bus, but may be connected by other interface and bus structures, such as a parallel port, game port or a universal serial bus (USB). A monitor 191 or other type of display device is also connected to the system bus 121 via an interface, such as a video interface 190. In addition to the monitor, computers may also include other peripheral output devices such as speakers 197 and printer 196, which may be connected through an output peripheral interface 195.

[0033] The computer 110 may operate in a networked environment using logical connections to one or more remote computers, such as a remote computer 180. The remote computer 180 may be a personal computer, a server, a router, a network PC, a peer device or other common network node, and typically includes many or all of the elements described above relative to the computer 110, although only a memory storage device 181 has been illustrated in FIG. 1. The logical connections depicted in FIG. 1 include a local area network (LAN) 171 and a wide area network (WAN) 173, but may also include other networks. Such networking environments are commonplace in offices, enterprise-wide computer networks, intranets and the Internet.

[0034] When used in a LAN networking environment, the computer 110 is connected to the LAN 171 through a network interface or adapter 170. When used in a WAN networking environment, the computer 110 typically includes a modem 172 or other means for establishing communications over the WAN 173, such as the Internet. The modem 172, which may be internal or external, may be connected to the system bus 121 via the user input interface 160, or other appropriate mechanism. In a networked environment, program modules depicted relative to the computer 110, or portions thereof, may be stored in the remote memory storage device. By way of example, and not limitation, FIG. 1 illustrates remote application programs 185 as residing on memory device 181. It will be appreciated that the network connections shown are exemplary and other means of establishing a communications link between the computers may be used.

Architecture

[0035] In one embodiment, the invention is implemented in or in conjunction with a web browser. A block diagram of an exemplary architecture is shown in Figure 2. The browser, 200, is a web browser which the user interacts with. A browser helper object 210 interfaces with the browser 200, captures browser events, and send the events to a user behavior tracer 220. The core

part of user behavior tracer 220 is state machine 225. The states of the state machine 225 represent the state of the user's interaction with browser 200. For example, one state may be associated with the user having initiated a search, but not yet received a result page. The state machine 225 receives events from user behavior tracer 220, tracks the state of the user's behavior as the user uses the browser 200, and collects implicit and explicit user feedback. The user behavior tracer 220 sends data regarding the search and the feedback to the data acquisition engine 230 which stores or processes the data.

[0036] In one embodiment, a Component Object Model (COM) or similar interface in browser 200 is used to allow an in-process COM server (or equivalent) to be developed which uses the interface in browser 200. This in-process COM server or equivalent is the browser helper object 210. The browser helper object 210 can, in one embodiment, be an in-process COM component registered under a certain registry key. When browser 200 is started, it looks up that registry key and loads the object whose class ID (CLSID) is stored there. (A CLSID is a globally unique identifier that identifies a COM class object.)

[0037] Browser helper object 210 detects the browser's typical events. In some browsers, events occur (are "fired") when certain situations occur. The detection of the event can signal that certain situations have occurred in the browser. For example, such events may include events which indicate the user's navigation from page to page, such as: (1) an event which fires when a hyper link has been clicked to navigate to a different page; (2) an event which fires when the history is used for navigation to a different page; (3) an event which fires when the address bar is used to navigate to a different page; and/or (4) an event which fires when the favorites list is used to navigate to a different page.

[0038] Additionally, such events may include events which indicate the user's use of a displayed page, such as: (5) an event which fires when a document has been completely loaded and initialized (6) an event which fires when scrolling is taking place; (7) an event which fires when a document is printed; (8) an event which fires when a document is added to the favorites list; (9) an event which fires when the window gains focus; (10) an event which fires when the window loses focus; and/or (11) an event which fires when a window has been closed. Other events may also exist and may be detected by the browser helper object 210. Some of these events may not be available in some browsers 200, and some browsers 200 may require modification in order to provide functionality for detection of these browsers.

[0039] In addition to notification of the occurrence of the corresponding activity in the browser, these events in some cases also include information about the activity. For example, when a hyperlink has been clicked (event 1) the element which has been clicked is sent as a parameter in the event.

[0040] In addition to detecting these events which signal the user's pattern of navigation and the user's use of a displayed page, browser helper object 210 can also access the browser's menu and toolbar and install hooks to monitor messages and actions which are not reflected in events. In one embodiment, these hooks are used instead of the use of events as detailed above.

[0041] Once this information is captured by the browser helper object 210, it is passed to the user behavior tracer 220. The events are used in two ways. First, the event may be stored as or used to calculate implicit feedback data. Such implicit feedback may include: data regarding the time spent on the page, mouse movements, mouse clicks, scrolling time, and query text entered. Additionally, implicit feedback may include event date and time stamps, result positions on page, link text, result URL, result click-through, result dwell-time, result printed, result added to My Favorites, scrolling activity detected on result, and whether the result was a sponsored link.

[0042] In addition to the use of the data from the browser helper object 210 for implicit feedback, the data, in one embodiment, is also used to transition from one state to another or to trigger an event in the state machine. Events in the state machine include user queries for feedback (such as feedback on a specific result or on a query in general) and user queries for state (for example, when a modified search may or may not be part of a new query, the user may be asked whether the user intended to continue the same query or start a new query.)

[0043] As stated above, in order to enable context-dependent evaluations of results of search mechanisms, context data is tracked. In one embodiment, this context data includes the state of the search, behavior data regarding the user's use of the browser 200 and, responses to user queries as to the intent of the user's behavior. Along with context data, user feedback data, either implicit or explicit, is also tracked. In this way, the user feedback can be used to evaluate the quality of the search mechanism in the context of the search performed. Additionally, relevance models can be evaluated as described above.

[0044] The state machine 225 works to detect the contours of a search session - when it starts, when it finishes, and what occurs during the search session. State machine 225 tracks the possible states of a search - when the user has finished the evaluation of the particular result, when the user is done with a particular query, and when questions need to be asked of users, such as what

feedback the user has on a particular result item or on the query as a whole. Similar events which are detected by the browser helper object 210 may mean different things, depending on what stage of the search session has been arrived at. In order to track this, in one embodiment, state machine 225 is used to handle the complexity of search events yet maintain flexibility. The state machine design pattern is a common pattern which can be implemented, for example, as detailed in the popular “Gang of Four” book entitled *Design Patterns: Elements of Reusable Object-Oriented Software* (Erich Gamma, Richard Helm, Ralph Johnson, and John Vlissides, Addison Wesley 1994.) Other implementations of state machine 225 are contemplated, as are other implementations of the user behavior tracer 220 which do not include a state machine.

State Machine

[0045] Figure 3 is a block diagram representing a state machine according to one embodiment of the invention. The state machine first detects that a search session has been started (for example, by browsing with browser 200 to a page with search capabilities and then entering a search). The state machine tracks the progress of one search, which may consist of one query or a number of queries. Each search corresponds to a search intent, which may be carried through more than one query. The determination of whether two queries are part of the same search, in one embodiment, is left to the user through explicit feedback. In other embodiments, this is determined implicitly, for example by the similarities between the user’s queries or the similarities in the results obtained.

[0046] When the data from the browser helper object 210 indicates that a search session is started, Search Session Start state 300 is entered. For example, this may occur when the user browses to a page with search mechanism capabilities. When and if data indicates that a query has been submitted (arrow 305), the Query Start state 310 is entered. The query submission to the search mechanism will result in a search result page being created at the search mechanism and downloaded by browser 200 (arrow 315). When this occurs, the state is transitioned to the Result List Page Visit state 320.

[0047] At this point, the user is looking at a page of results for the user’s query. From this state, Result List Page Visit state 320, the user may do different things. For example, if the user is looking at page N of M result pages from the search mechanism, the user may decide to go to a different result page. If this occurs, the state remains in Result List Page Visit state 320.

[0048] In one embodiment, the possibility of the user having finished the search is tracked and the user may be queried if it appears the query has ended. For example, the amount of time

spent in this state and/or the number of mouse clicks by the user is tracked, and if a click maximum threshold is exceeded or a timeout is encountered, the user is queried to determine if the user is still searching. If the user is not searching any more, a switch to an end state such as Query End state 360 or Search End state 370 is possible, as is an exit from the state machine, although these are not explicitly represented in Figure 3. At any state, the user may close the browser or open a page unrelated to the search, and a switch to such an end state or exiting of the state machine is possible.

[0049] If the search is not timed out or exited, though, the user can select a result item from the result list to examine. When this browsing to the item result (arrow 325) occurs, the Result Item Evaluation Start state 330 is entered. When the requested item result is downloaded to the browser 200 (arrow 335), the Result Item Page Visit state 340 is entered. In this state, the user is examining a result item. While the user browses the result item, for example, going from a first page to a second page, the state remains Result Item Page Visit state 340. While in this state, again a timeout may result in a query to the user and an exit to an end state.

[0050] The user may indicate that that the user is done visiting the result item explicitly, if a timeout occurs. Additionally, there are other ways in which the user's actions may indicate that the user is done visiting the result item. For example, the user may browse back to the result list, close the browser, or enter a new page in the address bar. When it has been determined that the user is done visiting the result item (arrow 345), the Evaluate Result Item state is entered. In this state, the user is queried for explicit feedback on the result item. For example, the user may be asked to rate the quality or usefulness of the result item on a 5 point scale by selecting a radio button on a popup window on the user's display.

[0051] When this evaluation is over, if the user has browsed back to the result list (arrow 353), the state returns to Result List Page Visit state 320. If the user has not, that is an indication that the query has ended (arrow 355) and the Query End state 360 is entered. In this state, in one embodiment, the user is queried for explicit feedback regarding the quality or usefulness of the query results as a whole. If the user indicates that the user wishes to refine the query (for example, either explicitly in response to a query)(arrow 363) the state returns to Query Start 310. Otherwise, the current search has ended (arrow 365) and the Search End state 370 is entered. In this state, in one embodiment, the user is queried for explicit feedback regarding the quality or usefulness of the search as a whole. From here the state machine may be exited or, if a new search is begun, the state machine may be reentered at Search Session Start state 300.

[0052] While transitions are occurring between the various states of the state machine 225 and explicit feedback is being requested and stored, implicit feedback and other user behavior information is being captured as well. For example, while the user is in the Result List Page Visit state 320, the users result list page explore behavior is being stored.

[0053] If there is a timeout, in one embodiment, a flag is set for the timeout. This is done because many computer systems allow users to use several applications at once. If the user has switched to another application, such as reading email with an email application, and the timeout occurs, the user experience may be negative if a dialog box pops up on the screen to ask him if he is done with the search. Instead, the flag is set. Then, when the user behavior tracer 220 receives data regarding an event indicating that the user is again using browser 200 (for example, by receiving notice that a mouse move event has been fired by the browser 200) then the user will be prompted for the timeout question.

[0054] Other state machines are also contemplated. For example, state machines which request different explicit user feedback are we grew the state machine by adding some more complicated transition cases. The state machine is able to handle quite complex scenarios, the complexity is well handled and controlled by the way we designed the state machine.

Data Recording

[0055] As discussed above, for each search, three types of data are acquired by the data acquisition engine 230 through the user behavior tracer 220.

[0056] (1) Context Data: This is data concerning the query or queries that the user used during the search. Context data can include data such as: the states of the search and associated timing, behavior data regarding the user's use of the browser 200 and responses to user queries as to the intent of the user's behavior. Generally context data allows user feedback data to be analyzed in the context of what was occurring leading to the feedback.

[0057] (2) Implicit User Feedback Data: This data reflects user behavior throughout the search session, such as page browsing, scrolling, clicking, and other user behavior. In one embodiment, timing information is also stored to describe the user behavior (As detailed above, as well as being used for implicit user feedback, data regarding user behavior is also used along with other data to drive transitions between the states in the state machine.) This user behavior data is not necessarily the raw user behavior data stream sent by the browser helper object 210. In one embodiment, the raw data is interpreted and stored as four types of implicit user feedback data:

- a) user behavior while visiting a result list page, including the time spent at that result;

- b) user behavior while exploring a hyper link on the result list page, such as a search tip link, including the time spent at that result;
- c) user behavior for visiting a result item page, including the time spent at that result and other actions such as scrolling, printing or adding the document to the user's favorites;

(in one embodiment, the active time spent at the result is measured in these three cases by subtracting any time in which the browser has lost user focus, for example when a user switches to another application)

- d) Result item ignore records the user behavior for ignoring a result item, i.e. when the user does not visit a result item though it appears as part of a result list.

[0058] These correspond to some extent to states in the state machine 225, and the interpretation, in one embodiment, of raw data into interpreted user behavior data of these types occurs in the states of the state machine, while they are being traversed.

[0059] (3) Explicit User Feedback Data: This data is the explicit user feedback which has been requested about the result items the user visited and the query the user have submitted. For example, regarding a search, a query, or a specific result, the user may be asked, via a dialog box, "Did this answer your question?" and allowed to enter a response. As another example, regarding a specific result which the user ignored, the user may be asked "Why didn't you try this result?" and given choices including "I didn't think this would answer my question."

[0060] In one embodiment, explicit user feedback is requested as shown in the following Table 1, with the description of the situation which will prompt the request for explicit feedback detailed in the Description column and the contents for the dialog interaction presented to the user as shown in the Dialog column with bullets representing radio buttons.

Description	Dialog
The user has finished evaluating a result item	Would you say that: <ul style="list-style-type: none"> • This result answered your question • This result somewhat answered your question • This result did not answer your question • You did not get a chance to evaluate this result (broken link, foreign language, etc.) • Other

The user exits a results list without exploring all results (the dialog will be presented about some or all of the results not visited, one dialog per result for which feedback is requested)	You did not visit this result because: <ul style="list-style-type: none"> • You didn't think it would answer your question • You've been to this result before • You didn't look at it • Other
A query has ended	For this query were you: <ul style="list-style-type: none"> • Satisfied with the query • Partially satisfied with the query • Not satisfied with the query
A search session has ended	For this search were you: <ul style="list-style-type: none"> • Satisfied with the search • Partially satisfied with the search • Not satisfied with the search

Table 1: User Feedback.

[0061] The query data, user behavior data, and explicit user feedback data which is collected can be presented for examination, used to validate a relevance model, or used to measure the effectiveness of a search mechanism.

[0062] Figure 4 is a flow diagram showing the process of gathering context-based feedback according to one embodiment of the invention. In step 400, the search mechanism is monitored for user behavior data regarding a search. At the same time, in step 410, the search mechanism is monitored for search mechanism response data regarding the search. This data is used to determine context data describing the search and to determine user feedback data describing the search, steps 420 and 430.

Conclusion

[0063] It is noted that the foregoing examples have been provided merely for the purpose of explanation and are in no way to be construed as limiting of the present invention. While the

invention has been described with reference to various embodiments, it is understood that the words which have been used herein are words of description and illustration, rather than words of limitations. Further, although the invention has been described herein with reference to particular means, materials and embodiments, the invention is not intended to be limited to the particulars disclosed herein; rather, the invention extends to all functionally equivalent structures, methods and uses, such as are within the scope of the appended claims. Those skilled in the art, having the benefit of the teachings of this specification, may effect numerous modifications thereto and changes may be made without departing from the scope and spirit of the invention in its aspects.